arXiv:1612.06675v1 [cs.DS] 20 Dec 2016

Clustering Uncertain Graphs

Matteo Ceccarello[1], Carlo Fantozzi[1], Andrea Pietracaprina[1],
Geppino Pucci[1] and Fabio Vandin[1]

[1] Department of Information Engineering, University of Padova, Italy,
{ceccarel,fantozzi,capri,geppo,vandinfa}@dei.unipd.it

Abstract

An uncertain graph $\mathcal{G} = (V, E, p : E \to (0, 1])$ can be viewed as a probability space whose outcomes (referred to as *possible worlds*) are subgraphs of $\mathcal{G}$ where any edge $e \in E$ occurs with probability $p(e)$, independently of the other edges. Uncertain graphs naturally arise in such diverse areas as computational biology, social network analysis, network reliability, and privacy enforcement, among the others, and very often their exploitation requires some form of probabilistic clustering of the node set. In this paper, we develop novel strategies for clustering uncertain graphs. Specifically, given an uncertain graph $\mathcal{G}$ and an integer τ, we aim at partitioning its nodes into τ clusters, each featuring a distinguished center node, so to maximize the smallest connection probability of any node to its cluster's center, in a random possible world. We prove that this problem is NP-hard, even in the presence of an oracle for the connection probabilities, and develop efficient approximation algorithms for the problem and for some useful variants. Unlike previous works in the literature, our algorithms feature provable approximation guarantees and are capable to keep the granularity of the returned clustering under control. Our theoretical findings are complemented by the results of several experiments which we run to compare our algorithms against the state-of-the-art competitor with respect to both running-time and quality of the returned clusterings.

1 Introduction

Currently, many application domains increasingly resort to the analysis of graphs that represent interrelated data. Oftentimes, the nature of the data is such that their associated graph features some level of uncertainty, which in some cases may be modeled by associating an existence probability to each link. For example, in Protein-Protein Interaction (PPI) networks, an edge between two proteins corresponds to an interaction that is observed through a noisy experiment characterized by some level of uncertainty, which can thus be conveniently cast as the probability of existence of that edge [1]. Also, in social networks, the probability of existence of an edge between two individuals may be used to model the likelihood of an interaction between the two individuals, or the influence of one of the two over the other [2]. Other applications of uncertainty in graphs arise in the realm of mobile ad-hoc networks [3, 4] and graph obfuscation for privacy enforcement [5].

The main objective of this work is to introduce novel strategies for clustering uncertain graphs, aiming at partitioning the node set so to maximize a connectivity-related measure defined in the next section. In accordance with the mainstream literature, an *uncertain graph* is defined over a set of nodes V, a set of edges E between nodes of V, and a probability function $p : E \to (0, 1]$. We denote such a graph by $\mathcal{G} = (V, E, p)$. $\mathcal{G}$ can be viewed as a probability space whose outcomes (referred to as *possible worlds*) are graphs $G = (V, E')$ where any edge $e \in E$ is included in E' with probability $p(e)$, independently of the other edges.

In the next subsection we provide a brief account on the most prominent works on uncertain graphs.

1.1 Related work

Early work on network reliability has dealt implicitly with the concept of uncertain graph. In general, given an uncertain graph, if we interpret edge probabilities as the complement of failure probabilities, a typical objective of network reliability analysis is to determine the probability that a given set of nodes is connected under random failures. This probability can be estimated through a Monte-Carlo approach

which becomes prohibitively cumbersome for very low reliability values. In fact, even the simplest problem of estimating the probability that two distinguished nodes s and t are connected is known to be $\#P$-complete [6, 7]. In the last three decades, several works have tried to come up with better heuristics for various reliability problems on uncertain graphs (see [8] and references therein).

The definition of uncertain graph adopted in this paper has been introduced in [9], where the authors investigate various probabilistic notions of distance between nodes, and develop efficient algorithm for determining the k nearest nodes of a given source under the various distance measures. It has to be remarked that the proposed measures do not satisfy the triangle inequality, thus ruling out the applicability of traditional graph clustering approaches [10]. In the last few years, there has been a multitude of works studying several analytic and mining problems on uncertain graphs. A detailed account of the state of the art on the subject can be found in [11], where the authors also investigate the problem of extracting a representative possible world providing a good summary of an uncertain graph for the purposes of query processing.

A number of recent works have studied different ways of clustering uncertain graphs, which is the focus of this paper. In [12] the authors consider, as a clustering problem, the identification of a deterministic *cluster graph*, which corresponds to a clique-cover of the nodes of the uncertain graph, aiming at minimizing the expected *edit distance* between the clique-cover and a random possible world of the uncertain graph, where the edit distance is measured in terms of edge additions and deletions. A 5-approximation algorithm for this problem is provided in the paper. The main drawback of this approach is that the formulation of the clustering problem does not allow to control the number of clusters. Moreover, the approximate solution returned by the proposed algorithm relies on a shallow star-decomposition of the topology of the uncertain graph, which always yields a large number of clusters (at least $V/(\Delta+1)$, where Δ is the maximum degree of a node in V). Thus, the returned clustering does not seem to exploit more systemic information about the connectivity properties of the underlying topology.

The same clustering problem considered in [12] has been also studied by Gu et al. in [13] for a more general class of uncertain graphs, where the assumption of edge independence is lifted and the existence of an edge (u, v) is correlated to the existence of its adjacent edges (i.e., edges incident on either u or v). The authors propose two algorithms for this problem, one that, as in [12], does not fix a bound on the number of returned clusters, and another that fixes such a bound. Neither algorithm is shown to provide worst-case guarantees on the approximation ratio.

In [14] a clustering problem is defined with the objective of minimizing the expected entropy of the returned clustering, defined with respect to the adherence of the clustering to the connected components of a random possible world. With this objective in mind, the authors develop a clustering algorithm which combines a standard k-means strategy with the Monte-Carlo sampling approach for reliability estimation. No theoretical guarantee is offered on the quality of the returned clustering with respect to the defined objective function, and the complexity of the approach does not appear to scale well with the size of the graph and also depends on a convergence parameter which cannot be estimated analytically. In summary, while the pursued approach to clustering has merit, there is no rigorous analysis of the tradeoffs that can be exercised between the quality of the returned clustering and the running time of the algorithm.

Boldi et al. [5] studied the problem of adding suitable edge probabilities to turn a deterministic graph G into an uncertain graph $\mathcal{G}$ for the purposes of identity obfuscation. The main goal of their work is to synthetize $\mathcal{G}$ in such a way that individual vertices of the original graph G can no longer be identified by their topological properties in $\mathcal{G}$, but at the same time $\mathcal{G}$ preserves (in expectation, over all possible worlds) some global properties of G, such as degree distribution, diameter, and clustering coefficient. In particular, the authors introduce the concept of a (k, ε)-obfuscation with respect to a given property P of the original graph G, where k relates to the entropy of the distribution of property P in the resulting uncertain graph $\mathcal{G}$, with higher values of k yielding higher levels of obfuscation, and $\varepsilon \in (0, 1]$ is a tolerance parameter, controlling the fraction of nodes which are not obfuscated. The authors provide an implementation of their obfuscation algorithm, which we will use in Section 5 to generate synthetic uncertain graphs upon which our clustering algorithms are exercised.

In [15] the *Markov Cluster Algorithm* (MCL) is proposed for clustering weighted graphs. In MCL, an edge weight is considered as a *similarity score* between the endpoints. The algorithm does not specifically target uncertain graphs, but it can be run on these graphs by considering the edge probabilities as weights. In fact, some of the aforementioned works on the clustering of uncertain graphs have used MCL for comparison purposes. The algorithm focuses on finding so-called *natural clusters*, that is, sets of nodes characterized by the presence of many edges and paths between their members. The basic idea of the

algorithm is to perform random walks on the graph and to partition the nodes into clusters according to the probability of a random walk to stay within a given cluster. Edge weights (i.e., the similarity scores) are used by the algorithm to define the probability that a given random walk traverses a given edge. The algorithm's behaviour is controlled with a parameter, called *inflation*, which indirectly controls the granularity of the clustering. However, there is no fixed analytic relation between the inflation parameter and the number of returned clusters, since the impact of the inflation parameter is heavily dependent on the graph's topology and the edge weights. The author of the algorithm maintains an optimized and very efficient implementation of MCL, against which we will compare our algorithms in Section 5.

1.2 Our Contribution

In this paper we develop novel strategies for clustering uncertain graphs. As observed in [14], a good clustering should aim at high probabilities of connectivity within clusters, which is however a hard goal to pursue, both because of the inherent difficulty of clustering per se, and because of the aforementioned #P completeness of reliability estimation in the specific uncertain graph scenario. Also, it has been observed [9, 14] that the straightforward reduction to the analysis of a deterministic graph where edge probabilities become weights may not yield significant outcomes because it disregards the possible world semantics.

Motivated by the above scenario, we will adopt the *connection probability* between two nodes (a.k.a. two-terminal reliability), that is, the probability that the two nodes belong to the same connected component in a random possible world, as the distance measure upon which we will base our clustering. As a first technical contribution, which may be of independent interest for the broader area of network reliability, we show that this measure satisfies some form of triangle inequality, unlike other distance measures used in previous works. This property allows us to cast the problem of clustering uncertain graphs into the same framework as traditional clustering approaches on metric spaces, while still enabling an effective integration with the possible world semantics that must be taken into account in the estimation of the connection probability.

Specifically, we study the following clustering problem, together with some variations. Given in input an uncertain graph $\mathcal{G}$ and an integer τ, we seek to partition the nodes of $\mathcal{G}$ into τ clusters, where each cluster contains a distinguished node, called *center*, such that the minimum connection probability of a node to its cluster's center is maximized. We first prove that the problem is NP-hard even in the presence of an oracle for the connection probability. Then, we develop a simple deterministic strategy which returns a τ-clustering where the minimum connection probability of a node to its cluster's center is $\Omega\left(p^2_{\text{opt}}(\tau)\right)$, where $p_{\text{opt}}(\tau)$ is the optimal connection probability of any τ-clustering. We also show how to adapt the strategy to approximate the somewhat dual problem of determining, for a fixed threshold q for the connection probability, a clustering of minimum cardinality whose featured connection probability is at least q.

Our strategy essentially grows the τ clusters sequentially one after the other, under decreasing guesses of the optimal connection probability. In order to speed up the cluster growing process (which is particularly useful for large graphs and large values of τ), we also develop an alternative strategy made of a number of parallel cluster growing phases which is logarithmic in $|V|$. This alternative strategy attains the same approximation quality of the sequential algorithm at the expense of an increase in the number of clusters by an $O\left(\log^2 |V|\right)$ factor. We finally present a variant of our algorithms that allows to impose a limit on the length of the paths that contribute to the connection probability between two nodes. Computing provably good clusterings under limited path length may have interesting applications in scenarios such as the analysis of PPI networks, where topological distance between two nodes diminishes their similarity, regardless of their connection probability.

In the paper, we first present our clustering strategies under the availability an oracle for the connection probabilities between pairs of nodes. Then, we show how to integrate a progressive sampling scheme for the Monte-Carlo estimation of the required probabilities which essentially preserves the approximation quality. While the Monte-Carlo computation of two-terminal reliability can be very computationally intensive for small probabilities, our clustering strategies tackle this obstacle by only requiring precise estimates of connection probabilities not much smaller than the max-min probability defining the optimal clustering quality.

We perform a preliminary experimental comparison of our algorithms with MCL on uncertain graphs derived from PPI networks and from DBLP. The results show that our algorithms identify good clusterings with respect to our metric, while a clustering strategy that is not specifically designed for uncertain graphs, as the one employed by MCL, may in some cases provide a very poor clustering. In addition,

for the biggest DBLP graph our algorithms are much faster than MCL. Moreover, the comparison of the clusterings of PPI networks obtained by our algorithms with a ground-truth shows that our algorithms are comparable with state of the art solutions in predicting members of the same protein complex. We also assess the relative performance of our sequential and parallel strategies.

The rest of the paper is organized as follows. In Section 2 we define basic concepts regarding uncertain graphs and the clustering problem under consideration. Also, we prove a form of triangle inequality for our connection probability measure and the NP-hardness of our clustering problem. In Section 3 we describe and analyze our clustering algorithms. In Section 4 we show how to integrate the Monte-Carlo probability estimation within the algorithms while maintaining comparable approximation guarantees. Section 5 reports the results of the experiments. Finally, Section 6 offers some concluding remarks and discusses possible avenues of future research.

2 Preliminaries

Let $\mathcal{G} = (V, E, p)$ be an uncertain graph, as defined in the introduction. In accordance with the established notation used in previous work, we write $G \sqsubseteq \mathcal{G}$ to denote that G is a possible world of $\mathcal{G}$. Given two nodes $u, v \in V$, the probability that they are connected (an event denoted as $u \sim v$) in a random possible world can be defined as

$$\Pr(u \sim v) = \sum_{G \sqsubseteq \mathcal{G}} \Pr(G)\mathbf{I}_G(u, v),$$

where

$$\mathbf{I}_G(u, v) = \begin{cases} 1 & \text{if } u \sim v \text{ in } G \\ 0 & \text{otherwise} \end{cases}$$

We refer to $\Pr(u \sim v)$ as the *connection probability* between u and v in $\mathcal{G}$. The uncertain graphs we consider in this paper, hence their possible worlds, are undirected and, except for the edge probabilities, no weights are attached to their nodes/edges.

Given an integer τ, with $1 \leq \tau < n$, a *τ-clustering* is a partition of V into τ *clusters* $C_1, \ldots, C_\tau$ and a set of *centers* $c_1, \ldots, c_\tau$ with $c_i \in C_i$, for $1 \leq i \leq \tau$. We study the problem of determining a τ-clustering which maximizes the following objective function

$$\min_{1 \leq i \leq \tau} \min_{v \in C_i} \Pr(c_i \sim v). \tag{1}$$

In other words, we want the clusters to be well connected, in the sense that each node is likely to be connected to its cluster's center in a random possible world. We can rephrase this problem as an instance of the well known k-center problem, which is known to be NP-hard [16], as follows. For every pair of nodes $u, v \in V$ define their distance as $d(u, v) = \ln(1/\Pr(u \sim v))$ with the understanding that $d(u, v) = \infty$ if $\Pr(u \sim v) = 0$. Under this transformation, the aim is the minimization of the following objective function

$$\max_{1 \leq i \leq \tau} \max_{v \in C_i} d(c_i, v). \tag{2}$$

When the distance function defines a metric space, and, in particular, it satisfies the triangle inequality (i.e, $d(u, z) \leq d(u, v) + d(v, z)$) the problem remains NP-hard but efficient approximation algorithms are known. It is easy to see that the triangle inequality for the above distance function can be expressed in terms of probabilities as $\Pr(u \sim z) \geq \Pr(u \sim v) \cdot \Pr(v \sim z)$, which we prove below. In contrast, the k-center problem cannot be immediately rephrased as the problem of finding a τ-clustering of a uncertain graph, since in the latter distances are constrained to arise from probabilities of a underlying uncertain graph.

Fix an arbitrary edge $e \in E$ and let $A(e)$ be the event: "edge e is present". We need the following technical lemma.

Lemma 1. *For any pair $x, y \in V$, we have $\Pr(x \sim y | A(e)) \geq \Pr(x \sim y | \neg A(e))$*

Proof. Let $\mathcal{G}_e^{x,y}$ (resp., $\mathcal{G}_{\neg e}^{x,y}$) be the set of possible worlds where $x \sim y$, and edge e is present (resp., not present). We have that

$$\begin{aligned} \Pr(x \sim y | A(e)) &= \sum_{G \sqsubseteq \mathcal{G}_e^{x,y}} \Pr(G)/p(e) \\ \Pr(x \sim y | \neg A(e)) &= \sum_{G \sqsubseteq \mathcal{G}_{\neg e}^{x,y}} \Pr(G)/(1 - p(e)). \end{aligned}$$

The lemma follows by observing that for any graph G in $\mathcal{G}^{x,y}_{\neg e}$ the same graph with the addition of e belongs to $\mathcal{G}^{x,y}_{e}$, and the corresponding terms in the two summations are equal. □

Theorem 1. *For any uncertain graph $\mathcal{G} = (V, E, p)$ and any triplet $u, v, z \in V$, we have:*

$$\Pr(u \sim z) \geq \Pr(u \sim v) \cdot \Pr(v \sim z).$$

Proof. The proof proceeds by induction on the number k of *uncertain edges*, that is, edges $e \in E$ with $p(e) \neq 0$ and $p(e) \neq 1$. Fix three nodes $u, v, z \in V$. The base case $k = 0$ is trivial, since, in this case, the uncertain graph is deterministic and for each pair of nodes $x, y \in V$, $\Pr(x \sim y)$ is either 1 or 0, which implies that when $\Pr(u \sim v) \cdot \Pr(v \sim z) = 1$, then $\Pr(u \sim z) = 1$ as well. Suppose that the property holds for uncertain graphs with at most k uncertain edges, with $k \geq 0$, and consider an uncertain graph $\mathcal{G} = (V, E, p)$ with $k + 1$ uncertain edges. Fix an arbitrary edge $e \in E$ and let A_e denote the event that edge e is present. For any two arbitrary nodes $x, y \in V$, we can write

$$\begin{aligned} \Pr(x \sim y) &= \Pr(x \sim y | A(e)) \cdot p(e) \\ &\quad + \Pr(x \sim y | \neg A(e)) \cdot (1 - p(e)) \\ &= (\Pr(x \sim y | A(e)) \\ &\quad - \Pr(x \sim y | \neg A(e))) \cdot p(e) + \Pr(x \sim y | \neg A(e)). \end{aligned} \tag{3}$$

By Lemma 1, the term multiplying $p(e)$ in the above expression is nonnegative. As a consequence, we have that

$$\Pr(u \sim v) \cdot \Pr(v \sim z) - \Pr(u \sim z) = A \cdot (p(e))^2 + B \cdot p(e) + C,$$

for some constants A, B, C independent of $p(e)$, with $A \geq 0$. Therefore, the maximum value of $\Pr(u \sim v) \cdot \Pr(v \sim z) - \Pr(u \sim z)$, as a function of $p(e)$, is attained for $p(e) = 0$ or $p(e) = 1$. Since in either case the number of uncertain edges is decremented by one, by the inductive hypothesis, the difference must yield a nonpositive value, hence the theorem follows. □

A fundamental primitive required for obtaining the desired clustering is the estimation of $\Pr(u \sim v)$ for any two nodes $u, v \in V$. While the exact computation of $\Pr(u \sim v)$ is $\#P$-complete [7], for reasonably large values of this probability a very accurate estimate can be obtained through Monte-Carlo sampling. More precisely, for $r > 0$ let $G_1, \ldots, G_r$ be r sample possible worlds drawn independently at random from $\mathcal{G}$. For any pair of nodes u and v we can define the following estimator

$$\tilde{q}(u, v) = \frac{1}{r} \sum_{i=1}^{r} \mathbf{I}_{G_i}(u, v) \tag{4}$$

It is easy to see that $\tilde{q}(u, v)$ is an unbiased estimator of $\Pr(u \sim v)$. Moreover, by taking

$$r \geq \frac{3 \ln \frac{2}{\delta}}{\varepsilon^2 \Pr(u \sim v)} \tag{5}$$

samples, we have that $\tilde{q}(u, v)$ is an (ε, δ)-approximation of $\Pr(u \sim v)$, that is,

$$\Pr\left(\frac{|\tilde{q}(u, v) - \Pr(u \sim v)|}{\Pr(u \sim v)} \leq \varepsilon \right) \geq 1 - \delta \tag{6}$$

(e.g., see [17, Theorem 10.1]). This approach is very effective when $\Pr(u \sim v)$ is not too small. However, when $\Pr(u \sim v)$ approaches 0, the number of samples, hence the work, required to attain an accurate estimation becomes prohibitively large.

Even if probabilities $\Pr(u \sim v)$ were provided by an oracle (i.e., could be computed efficiently), the problem would still be computationally difficult, as we now prove. We consider the following decision problem: given an uncertain graph $\mathcal{G} = (V, E, p)$, an integer $\tau \geq 1$, and $\hat{p}$ with $0 \leq \hat{p} \leq 1$, is there a τ-clustering such that $\min_{1 \leq i \leq \tau} \min_{v \in C_i} \Pr(c_i \sim v) \geq \hat{p}$?

Theorem 2. *The decision problem above is NP-hard.*

(sketch). The proof is by reduction from set cover. Consider an instance of set cover, with $U = \{u_1, \ldots, u_m\}$ denoting the universe of elements and $\mathcal{S} = \{S_1, \ldots, S_n\}$, with $S_i \subseteq U, \forall 1 \leq i \leq n$, being the elements *covered* by the sets. Given an integer k, the set cover problem asks whether there are k sets that cover all the elements of U, that is $S_{i_1}, S_{i_2}, \ldots, S_{i_k}$ with $S_{i_j} \in \mathcal{S}$ $\forall 1 \leq j \leq k$, such that $\cup_{j=1}^{k} S_{i_j} = U$. Note that if there exists an element $u \in U$ such that $u \notin S_i$ for $1 \leq i \leq n$ there cannot be a set cover for U. Since such condition can be checked in polynomial time, we now assume that such condition does not hold.

Given an instance of set cover, we build an instance of our decision problem as follows. Let $N = m+n$. We consider the uncertain graph $\mathcal{G} = (V, E, p)$ where: $V = U \cup \mathcal{S}$; $E = \{(u, S) : S \in \mathcal{S}, u \in S\} \cup \{(S, S') : S, S' \in \mathcal{S}\}$; $p(e) = \frac{1}{N!}$. Note that this procedure takes time polynomial in mn.

We now prove that $\mathcal{G} = (V, E, p)$ admits a k-clustering such that $\min_{1 \leq i \leq \tau} \min_{v \in C_i} \Pr(c_i \sim v) \geq \frac{1}{N!}$ if and only if there is a set cover of cardinality k. We first prove that if there is a set cover of cardinality k, then there is a clustering with minimum probability $\geq \frac{1}{N!}$. Define $S_{i_1}, S_{i_2}, \ldots, S_{i_k}$ to be centers of the k clusters, and assign each vertex u in U to a cluster of center S_{i_j} with $u \in S_{i_j}$ (note that such center must exist since $S_{i_1}, S_{i_2}, \ldots, S_{i_k}$ is a set cover); each vertex $S \in \mathcal{S} \setminus \{S_{i_1}, S_{i_2}, \ldots, S_{i_k}\}$ is assigned to an arbitrary cluster. Note that by construction, for each vertex $u \in V \setminus \{S_{i_1}, S_{i_2}, \ldots, S_{i_k}\}$ and its corresponding center S_{i_j} in the clustering, $(u, S_{i_j}) \in E$; therefore $\Pr(u \sim S_{i_j}) \geq \frac{1}{N!}$, that is the minimum connection probability for the clustering is $\geq \frac{1}{N!}$.

We now prove that if there is a k clustering with minimum connection probability $\geq \frac{1}{N!}$ then there is a k set cover. We first prove that in every k clustering with minimum connection probability $\geq \frac{1}{N!}$ a vertex v of a cluster is directly connected to the center c of its cluster, that is $(u, c) \in E$. Let assume that this was not the case, and consider the node u and the center c of its cluster such that $(u, c) \notin E$. Note that the event "u is connected to c" is equivalent to "$\cup_{p \in SP}$ p exists", where SP is the set of all simple paths among u and c in $\mathcal{G}$. Thus, by union bound $\Pr(u \sim c) \leq \sum_{p \in SP} \Pr(p)$, with $\Pr(p)$ being the probability that path p exists in a realization of $\mathcal{G}$. Note that there are $< N$ paths of length 2 in SP each with $\Pr(p) = \left(\frac{1}{N!}\right)^2$, and $< N!N$ paths of length > 2 in SP, each with probability $\Pr(p) \leq \left(\frac{1}{N!}\right)^3$. Thefore if u and c are not directly connected $\Pr(u \sim c) < \frac{1}{N!}$, that is a contradiction. If all the centers of the clustering are in the set $\mathcal{S}$, then the centers cover the universe U of elements, since for every vertex $u \in U$ the center S of the cluster u belongs to contains u. Otherwise, we can transform the clustering in a clustering with all centers in $\mathcal{S}$ and the same minimum connection probability: consider a center $c \in U$ and let $C \subset V \setminus \{c\}$ be the set of vertices other than c in the cluster with center c. If $C = \emptyset$, reassign the center of cluster $C \cup \{c\}$ to an arbitrarily chosen element of $\mathcal{S}$ that covers c. (If all elements of $\mathcal{S}$ covering c are already centers, assign c to an arbitrary cluster among the ones with such centers.) Otherwise, note that C contains only vertices of $\mathcal{S}$, since there is no edge among vertices in U. Now reassign the center of cluster $C \cup \{c\}$ to an arbitrary vertex of C; note that since all vertices in C are connected by an edge, the minimum connection probability for the new clustering is $\geq \frac{1}{N!}$. □

3 Clustering algorithms

A natural approach for finding a τ-clustering of an uncertain graph would consist of adapting the well-known approximation algorithm devised in [18] for clustering a given pointset around suitably selected centers so to minimize the distance of each point from its cluster's center. However, this algorithm relies on the repeated determination of the *farthest* point from a subset of currently selected centers and, when adapted to the uncertain graph scenario, it might entail the computation of very small connection probabilities between pairs of nodes, which is computationally hard, as we pointed out in the previous section. As a consequence, we have to resort to strategies which are able to create clusters based on *proximity* rather than on *farness*. In this section, we present a number of proximity-based approximation algorithms for clustering uncertain graphs, with provable approximation guarantees. First, in Subsection 3.1 we describe a simple deterministic algorithm for the τ-clustering problem introduced before, which is based on a sequential cluster growth strategy, and a variant of this algorithm that, for given threshold for the connection probability, computes a clustering of small granularity where each node is connected to its cluster's center with probability not smaller than the threshold. Then, in Subsection 3.2 we describe a randomized algorithm that uses a more efficient parallel cluster growing strategy while still providing good guarantees on the achieved connection probability, at the expense of a slight (polylogarithmic) increase in the number of clusters. Finally, in Subsection 3.3 we show how the approximation guarantees of the proposed algorithms change when an upper bound on the length of the connecting paths between

Algorithm 1: SEQUENTIALCLUSTERING1$(\mathcal{G}, \tau, \gamma)$

```
p_curr ← 1;
while true do
    S ← ∅;                                                    ▷ Set of centers
    V_1 ← V;
    for i ← 1 to τ do
        select one arbitrary cluster center c_i ∈ V_i;
        M_i ← {u ∈ V_i : Pr(c_i ∼ u) ≥ p_curr};
        V_{i+1} ← V_i − M_i;
        S ← S ∪ {c_i};
        if V_{i+1} = ∅ then
            for j ← 1 to i do
                compute C_j = {u ∈ V : c(u, S) = c_j};
            end
            return clusters C_1, C_2, ..., C_i centered at    c_1, c_2, ..., c_i;
        end
    end
    p_curr ← p_curr/(1 + γ)
end
```

the center and the other nodes of a cluster is imposed.

All algorithms presented in this section take as input an uncertain graph $\mathcal{G} = (V, E, p)$ with n nodes, and assume the existence of an oracle that given two nodes $u, v \in V$ returns $\Pr(u \sim v)$. In Section 4 we will discuss how to integrate the Monte-Carlo estimation of such probability within our algorithms.

3.1 Sequential cluster growing

For a given target number τ of clusters, with $1 \leq \tau \leq n$, let us consider an optimal partition of V into τ clusters $\hat{C}_1, \ldots, \hat{C}_\tau$, with centers $\hat{c}_1, \ldots, \hat{c}_\tau$, which maximizes the objective function defined by Equation (1), and let $p_{\text{opt}}(\tau)$ denote the value of this objective function, namely,

$$p_{\text{opt}}(\tau) = \min_{1 \leq i \leq \tau} \min_{v \in \hat{C}_i} \Pr(\hat{c}_i \sim v).$$

Observe that $p_{\text{opt}}(\tau) > 0$ if and only if the uncertain graph $\mathcal{G}$ features no more than τ connected components. For convenience, we assume that this is the case. For a node $u \in V$ and a subset of centers $S \subset V$, let $c(u, S) = \arg\max_{c \in S}\{\Pr(c \sim u)\}$ denote the node in S connected to u with the highest probability (ties broken arbitrarily), and let $\pi(u, S) = \Pr(c(u, S) \sim u)$.

Algorithm 1 computes a τ-clustering of $\mathcal{G}$ using a simple sequential cluster growing strategy where a design parameter $\gamma \in (0, 1)$ can be exploited to exercise a quality-performance tradeoff. The algorithm maintains a threshold probability p_{curr} and proceeds iteratively. In Iteration i, it starts by selecting an arbitrary new center c_i among all those nodes which are not marked as connected, i.e. are not connected within probability p_{curr} to any of the previously selected centers. This operation is repeated until there are no more nodes to select as centers, or τ centers have already been selected. In the former case, the algorithm assigns each node $u \in V$ to the cluster centered in $c(u, S)$. In the latter, if there are some nodes not connected within probability p_{curr} to one of the τ selected centers, the algorithm discards all centers, decreases p_{curr} as a function of the parameter γ, and starts building a new clustering from scratch.

It is easy to see that the algorithm terminates returning an ℓ-clustering, with $\ell \leq \tau$, since p_{curr} cannot drop by a factor $1/(1 + \gamma)$ below the minimum connection probability between two nodes in the same connected component of $\mathcal{G}$ before the termination condition is reached. Let $C_1, \ldots, C_\ell$ be the clusters returned by the algorithm centered at $c_1, \ldots, c_\ell$, and let p_{alg} be the value of the objective function corresponding to this clustering, namely

$$p_{\text{alg}} = \min_{1 \leq i \leq \ell} \min_{v \in C_i} \Pr(c_i \sim v).$$

We have:

Algorithm 2: SEQUENTIALCLUSTERING2$(\mathcal{G}, q)$

$V_1 \leftarrow V$; $i \leftarrow 1$;
while *true* **do**
  select one arbitrary cluster center $c_i \in V_i$;
  $C_i \leftarrow \{u \in V_i \,:\, \Pr(c_i \sim u) \geq q\}$;
  $V_{i+1} \leftarrow V_i - C_i$;
  if $V_{i+1} = \emptyset$ **then return** clusters $C_1, C_2, \ldots, C_i$ centered at $c_1, c_2, \ldots, c_i$;
  $i \leftarrow i + 1$;
end

Theorem 3. *Algorithm 1 terminates after at most* $\lfloor 2\log_{1+\gamma}(1/p_{\rm opt}(\tau))\rfloor + 1$ *iterations of the while loop and returns a clustering with*

$$p_{\rm alg} \geq \frac{p^2_{\rm opt}(\tau)}{(1+\gamma)}$$

Proof. Let $\ell = \lfloor 2\log_{1+\gamma}(1/p_{\rm opt}(\tau))\rfloor + 1$. Suppose that the algorithm reaches the ℓ-th iteration of the while loop and note that at the beginning of this iteration $p^2_{\rm opt}(\tau)/(1+\gamma) < p_{\rm curr} \leq p^2_{\rm opt}(\tau)$. Let $\hat{C}_{k_i}$ be the cluster in the optimal τ-clustering which contains c_i (the center selected in the i-th iteration of the for loop), for every $i \geq 1$. Moreover, let $\hat{c}_{k_i}$ be the center of $\hat{C}_{k_i}$. By Theorem 1 we have that for every node $v \in \hat{C}_{k_i}$

$$\Pr(c_i \sim v) \geq \Pr(c_i \sim \hat{c}_{k_i}) \cdot \Pr(\hat{c}_{k_i} \sim v) \geq p^2_{\rm opt}(\tau) \geq p_{\rm curr}$$

Therefore, V_{i+1} cannot contain any node of $\hat{C}_{k_i}$. In fact, inductively, we have that V_{i+1} cannot contain any node of $\bigcup_{j=1}^{i} \hat{C}_{k_j}$. Therefore, V_{t+1} for some $t \leq \tau$ must be empty. The theorem follows by observing that $p_{\rm alg}$ cannot be smaller than the value $p_{\rm curr}$ used the last iteration of the while loop. □

In some application scenarios, it may be useful to control the connectivity within the clusters by imposing a lower bound on the connection probability of a node to the center of its cluster. Thus, we can define an alterative clustering problem which, given a probability threshold q, requires to find a partition of nodes into a minimum number of clusters where the connection probabilities of the nodes to their assigned clusters' centers is at least q. Let $\tau(q)$ denote such minimum number of clusters. It is easy to see that $\tau(q)$ is a non decreasing function of q. Algorithm 2 computes an approximate solution to this problem. The algorithm proceeds iteratively: in iteration i, it selects from the unassigned nodes an arbitrary vertex c_i as a new cluster center and greedily places all unassigned nodes $v \in V$ with $\Pr(c_i \sim v \geq q)$ to the cluster centered at c_i. We have:

Theorem 4. *Algorithm 2 returns a* τ*-clustering with* $\tau_{alg} \leq \tau_{opt}(\sqrt{q})$*, where each node is connected to the center of its cluster with probability at least* q*.*

Proof. By construction, the algorithm generates clusters featuring the given connection probability. Consider an optimal $\tau(\sqrt{q})$-clustering, where each node is connected to the center of its cluster with probability at least $\sqrt{q}$. Note that by Theorem 1, two nodes c, u belonging to the same cluster of this optimal $\tau(\sqrt{q})$-clustering are connected with probability $\geq q$. An argument similar to the one used in the proof of Theorem 3 shows that at the end of the i-th iteration of the while loop all nodes which in the optimal $\tau(\sqrt{q})$-clustering belong to the same cluster as one of the c_j's, with $j \leq i$, must have been already included in some cluster. Hence, no more than $\tau_{\rm opt}(\sqrt{q})$ clusters are returned by the algorithm. □

3.2 Parallel cluster growing

There are two shortcomings of the τ-clustering algorithm presented in the previous subsection (Algorithm 1). First, for each of the guesses of the connection probability, τ clusters must be computed sequentially one after the other, which may be computationally cumbersome for large graphs and large values of τ. Second, the clustering is rigid with respect to the upper bound τ on the number of clusters. In fact, allowing for a slightly larger number of clusters might provide, in some cases, a solution of better quality. In this subsection, we propose an alternative clustering algorithm (Algorithm 3) inspired by the strategy developed in [19] for deterministic graphs, which grows many clusters in parallel, featuring at the same time an implicit adaptation of the connection probability.

Algorithm 3: ParallelClustering$(\mathcal{G}, \tau)$

$V_1 \leftarrow V$; $i \leftarrow 1$; $S \leftarrow \emptyset$;
while $|V_i| > \alpha\tau \log n$ **do**
  $S_i \leftarrow \{\alpha\tau \log n$ random centers from $V_i\}$;
  let $u_1, u_2, \ldots, u_{|V_i|}$ be the nodes of V_i in decreasing order of $\pi(\cdot, S_i)$;
  $S \leftarrow S \cup S_i$;
  $p_i \leftarrow \pi(u_{|V_i|/2}, S_i)$;
  foreach $u \in V_i$ **do**
    if $\pi(u, S) \geq p_i$ **then** assign u to cluster centered at $c(u, S)$;
  end
  $V_{i+1} \leftarrow \{u \in V_i \ : \ \pi(u, S) < p_i\}$;
  $i \leftarrow i + 1$;
end
return the clusters centered at $S \cup V_i$;

The algorithm is randomized and proceeds in *phases*, adding a new batch of cluster centers in each phase so to cover at least half of the yet uncovered nodes of the graph. The idea underlying this strategy is to push more clusters in regions of the graph with low connection probability: the clusters covering these regions will have fewer nodes and might thus avoid including nodes reachable only with very low probabilities. Let $\alpha > 0$ be a suitable constant which will be fixed by the analysis. Algorithm 3 works as follows. In Phase $i \geq 1$, it first selects at random a batch of $\alpha\tau \log n$ centers from the nodes that are still uncovered[1]. Let S_i be the set of newly selected centers, and let V_i be the set of yet uncovered nodes at the beginning of Phase i. The algorithm sorts the nodes $u \in V_i$ by decreasing order of $\pi(u, S_i)$, i.e., by decreasing probability of connection to some center of S_i. Let p_i be the connection probability of the $(|V_i|/2)$-th node in the ordering. Then, each node in V_i featuring a connection probability at least as large as p_i to one of the centers selected either in the current or in previous phases, is assigned to the most convenient cluster. When the number of unassigned nodes is reduced to at most $\alpha\tau \log n$, the algorithm terminates by considering all the "leftover" nodes as additional centers.

Suppose that the algorithm returns ℓ clusters, $C_1, \ldots, C_\ell$, centered at $\{c_1, \ldots, c_\ell\}$, and note that the clusters form indeed a partition of V. Let p_{alg} be the value of the objective function corresponding to this clustering, namely

$$p_{\text{alg}} = \min_{1 \leq i \leq \ell} \min_{v \in C_i} \Pr(c_i \sim v).$$

As before, we define $p_{\text{opt}}(\tau)$ as the objective function of an optimal τ-clustering $\hat{C}_1, \ldots, \hat{C}_\tau$, with centers $\hat{c}_1, \ldots, \hat{c}_\tau$. We have:

Theorem 5. *Algorithm 3 returns an ℓ-clustering with $\ell = O\left(\tau \log^2 n\right)$ and such that $p_{alg} \geq p^2_{opt}(\tau)$, with high probability.*

Proof. Let n_i be the number of nodes not yet assigned to clusters at the beginning of Phase i of the while loop (i.e., $n_i = |V_i|$), for $i \geq 1$. Consider an arbitrary Phase i and observe that, by construction, at least $n_i/2$ nodes of V_i are assigned to clusters during the phase. This immediately implies that at most $\log n$ phases are executed overall, and the bound on ℓ follows. In order to prove the lower bound on p_{alg} we show that $p_i \geq p^2_{\text{opt}}(\tau)$. To this purpose, it is sufficient to show that there exist at least $n_i/2$ nodes $u \in V_i$ such that $\pi(u, S) \geq \pi(u, S_i) \geq p^2_{\text{opt}}(\tau)$. Define an optimal cluster $\hat{C}_j$ to be *large* if $|\hat{C}_j \cap V_i| \geq (n_i/(2\tau))$. It is easy to see that the optimal clusters which are not large account for less than half of the nodes of V_i, hence large clusters contain at least $n_i/2$ nodes of V_i. We now argue that at least one center of S_i falls into every large cluster $\hat{C}_j$, with high probability. Consider an arbitrary large cluster $\hat{C}_j$. The probability that no center of S_i is selected among the nodes of $\hat{C}_j \cap V_i$ is

$$\left(1 - \frac{|\hat{C}_j \cap V_i|}{n_i}\right)^{\alpha\tau \log n} \leq \left(1 - \frac{1}{2\tau}\right)^{\alpha\tau \log n} \leq e^{-(1/2)\alpha \log n},$$

which can be made polynomially small in n for large enough α. Recall that $\hat{C}_j$ includes a node $\hat{c}_j$ such that $\Pr(\hat{c}_j \sim u) \geq p_{\text{opt}}(\tau)$, for every $u \in \hat{C}_j$. Suppose that a center $c \in \hat{C}_j \cap V_i$ is selected as a new center

[1] Unless otherwise specified, all logarithms in the paper are to base 2, while ln identifies the natural logarithm.

in the phase. Then, for every $u \in \hat{C}_j \cap V_i$, by Theorem 1 we have that

$$\Pr(c \sim u) \geq \Pr(c \sim \hat{c}_j) \cdot \Pr(\hat{c}_j \sim u) \geq p^2_{\text{opt}}(\tau).$$

This immediately implies that $\pi(u, S) \geq \Pr(c \sim u) \geq p^2_{\text{opt}}(\tau)$. Since there are at least $n_i/2$ nodes of V_i in large clusters, we can conclude that $p_i \geq p^2_{\text{opt}}(\tau)$. □

We remark that in contrast to Algorithm 1, Algorithm 3 lends itself to a more effective parallelization, since it requires only $O(\log n)$ parallel cluster growing steps, and attains a slightly better connection probability. However, these better features are counter-balanced by a larger number of clusters which, in the worst case, can be a factor $O\left(\log^2 n\right)$ larger than the target τ. However, in Section 5 we will provide some evidence that the actual number of returned clusters is much less than the worst case.

3.3 Limiting the path length

The algorithms described in the preceding subsections can be run by setting a limit on the length of the paths that contribute to the connection probability between two nodes. More precisely, for a fixed integer d, with $1 \leq d < n$, we define $\Pr(u \overset{d}{\sim} v) = \sum_{G \sqsubseteq \mathcal{G}} \Pr(G)\mathbf{I}_G(u, v; d)$, where $\mathbf{I}_G(u, v; d)$ is 1 if u is at distance at most d from v in G, and 0 otherwise. In the following, we refer to $\Pr(u \overset{d}{\sim} v)$ as the *d-connection probability* between u and v. By easily adapting the proofs of Lemma 1 and Theorem 1, it can be shown that for any triplet $u, v, z \in V$ it holds

$$\Pr(u \overset{d}{\sim} z) \geq \Pr(u \overset{\lfloor d/2 \rfloor}{\sim} v) \cdot \Pr(v \overset{\lfloor d/2 \rfloor}{\sim} z).$$

For a τ-clustering $C_1, \ldots, C_\tau$ with centers $\{c_1, \ldots, c_\tau\}$, consider the following objective function

$$\min_{1 \leq i \leq \tau} \min_{v \in C_i} \Pr(c_i \overset{d}{\sim} v). \tag{7}$$

and let $p_{\text{opt}}(\tau, d)$ be the maximum value of this objective function exhibited by any τ-clustering. Observe that $p_{\text{opt}}(\tau, d) > 0$ if and only if there exists a τ-clustering of the uncertain graph where there exists a path of length at most d between each node and its assigned center. Suppose that we run Algorithms 1 and 3 using the d-connection probability rather than the unconstrained connection probability, anywhere it is required, and let p_{alg} denote the value of above objective function for the clusterings returned by the algorithms. We have the following theorem for the algorithm based on the sequential cluster growing strategy.

Theorem 6. *Suppose that $p_{opt}(\tau, \lfloor d/2 \rfloor) > 0$. Considering only d-connection probabilities, Algorithm 1 terminates after at most $\lfloor 2 \log_{1+\gamma}(1/p_{\text{opt}}(\tau, \lfloor d/2 \rfloor)) \rfloor + 1$ iterations of the while loop and returns a clustering with*

$$p_{alg} \geq \frac{p^2_{opt}(\tau, \lfloor d/2 \rfloor)}{(1+\gamma)}.$$

The assumption on $p_{\text{opt}}(\tau, \lfloor d/2 \rfloor) > 0$ is required in the above statement to ensure that the algorithm terminates. As for the algorithm based on the parallel cluster growing strategy, we have the following result.

Theorem 7. *Considering only d-connection probabilities, Algorithm 3 returns an ℓ-clustering with $\ell = O\left(\tau \log^2 n\right)$ and such that $p_{alg} \geq p^2_{opt}(\tau, \lfloor d/2 \rfloor)$, with high probability.*

The proofs of the two theorems above follow easily, with virtually the same arguments, by substituting $p_{\text{opt}}(\tau, \lfloor d/2 \rfloor)$ for $p_{\text{opt}}(\tau)$ in the proofs of Theorem 3 and 5, respectively.

As for Algorithm 2, again suppose we run it using the d-connection probability rather than the unconstrained connection probability, anywhere it is required. Define $\tau(\sqrt{q}, d)$ to be the minimum number of clusters that yield a value of the objective function of Equation 7 greater than or equal to the threshold q, and let τ_{alg} be the number of clusters returned by the algorithm. Then, by substituting $\tau(\sqrt{q})$ with $\tau(\sqrt{q}, \lfloor d/2 \rfloor)$ in the statement of Theorem 4, we can prove the following theorem with virtually the same arguments.

Theorem 8. *Considering only d-connection probabilities, Algorithm 2 returns a τ-clustering with $\tau_{alg} \leq \tau_{opt}(\sqrt{q}, \lfloor d/2 \rfloor)$, where each node is connected to the center of its cluster with probability at least q.*

4 Implementing the oracle

In the previous section we assumed that the probabilities $\Pr(u \sim v)$ could be obtained exactly from an oracle. In practice, the estimation of these probabilities becomes the most critical part for the efficient implementation of our algorithms. In this section, we show how to integrate the Monte-Carlo sampling method for the estimation of the connection probabilities within the algorithms described in Section 3, maintaining similar guarantees on the quality of the returned clusterings. The basic idea of our approach is to adjust the number of samples dynamically during the execution of the algorithms, based on safe guesses of the probabilities that need to be estimated. We show how to apply this idea to implement the sequential and parallel cluster growing strategies in Subsections 4.1 and 4.2, respectively.

Throughout the section, we assume that a lower bound p_L to $p^2_{\text{opt}}(\tau)$ is available. In theory, this lower bound can be obtained by squaring the probability of the most unlikely world, but in practice it may represent a minimum connection probability set by the user, which may not be interested in clusterings with very low connection probabilities. In this latter case, if the algorithm does not find a clustering using only probabilities above p_L, it terminates by reporting that no clustering has been found. Recall that $\tilde{q}(u, v)$ denotes the estimate of the probability $\Pr(u \sim v)$ obtained by sampling possible worlds (Equation 4). Moreover, for a node $u \in V$ and a set of nodes $S \subset V$, we define $\tilde{c}(u, S) = \arg\max_{c \in S}\{\tilde{q}(c, u)\}$ as the function returning the node of S connected to u with the highest *estimated* probability. Similarly, we define $\tilde{\pi}(u, S) = \Pr(\tilde{c}(u, S) \sim u) = \max_{c \in S}\{\tilde{q}(c, u)\}$. We use $\varepsilon > 0$ to denote an approximation parameter to be fixed by the user.

4.1 Sequential growing strategies

We first consider Algorithm 1 and then turn to Algorithm 2. Consider Algorithm 1, and recall that it uses a probability threshold p_{curr} which is lowered at each iteration. Algorithm 4 provides an implementation of Algorithm 1 which maintains the same overall structure and uses progressive sampling to estimate connection probabilities.

In any iteration of the while loop, Algorithm 4 seeks only connection probabilities that are $\geq p_{\text{curr}}$. Therefore, we can limit the number of samples used in the iteration just to the ones necessary for the reliable estimation of probabilities greater than p_{curr}. This number of samples is given by the following equation

$$r = \frac{12}{p_{\text{curr}}\varepsilon^2} \ln\left(\frac{2n^2}{\delta}\left(1 + \left\lfloor \log_{1+\gamma} \frac{1}{p_L} \right\rfloor\right)\right) \tag{8}$$

where $\delta \in O\left(n^{-d}\right)$ for some constant $d > 1$. To take into account the error introduced by the estimation, which is bounded by $\varepsilon/2$, we modify the rule with which nodes are marked as covered by cluster centers as follows: given the value of p_{curr}, the algorithm marks as covered by a center c all nodes u with estimated probability $\tilde{q}(c, u) \geq (1 - \frac{\varepsilon}{2})p_{curr}$.

Let $\tilde{p}_{\text{alg}}$ be the objective function of the clustering obtained from the progressive sampling implementation of the sequential clustering algorithm.

Theorem 9. *Suppose that $p^2_{opt}(\tau) > p_L$. Algorithm 4 terminates after at most $\lfloor 2\log_{1+\gamma}(1/p_{\text{opt}}(\tau))\rfloor + 1$ iterations of the while loop and returns a clustering with*

$$\tilde{p}_{alg} \geq \frac{(1-\varepsilon)}{(1+\gamma)} p^2_{opt}(\tau),$$

with high probability.

Proof. Consider an arbitrary iteration of the while loop, and a pair of nodes $u, v \in V$. Let

$$\delta' = \frac{\delta}{n^2\left(1 + \left\lfloor \log_{1+\gamma} \frac{1}{p_L} \right\rfloor\right)}$$

By the choice of the number of samples (Equation (8)) and Inequality (6) we have the following properties on the estimate $\tilde{q}(u, v)$

- if $\Pr(u \sim v) > p_{\text{curr}}$, then $\tilde{q}(u, v) < (1 - \frac{\varepsilon}{2})p_{\text{curr}}$ with probability $< \delta'$
- if $\Pr(u \sim v) < (1 - \varepsilon)p_{\text{curr}}$, then $\tilde{q}(u, v) \geq (1 - \frac{\varepsilon}{2})p_{\text{curr}}$ with probability $< \delta'$

Algorithm 4: SEQUENTIALCLUSTERIMPL($\mathcal{G}, \tau, \gamma$)

$p_{\text{curr}} \leftarrow 1$;
while $p_{curr} \geq p_L$ **do**
  $r \leftarrow \frac{12}{p_{\text{curr}}\varepsilon^2} \ln\left(\frac{2n^2}{\delta}\left(1 + \left\lfloor \log_{1+\gamma} \frac{1}{p_L} \right\rfloor\right)\right)$;
  $V_1 \leftarrow V$;
  $S \leftarrow \emptyset$; ▷ Set of centers
  Sample r possible worlds $G_1, \ldots, G_r \sqsubseteq \mathcal{G}$;
  for $i \leftarrow 1$ **to** τ **do**
    select one arbitrary cluster center $c_i \in V_i$;
    for $u \in V$ **do**
      $\tilde{q}(c_i, u) \leftarrow \frac{1}{r} \sum_{h=1}^{r} \mathbf{I}_{G_h}(u, v)$;
    end
    $M_i \leftarrow \{u \in V_i \ : \ \tilde{q}(c_i, u) \geq (1 - \frac{\varepsilon}{2}) p_{curr}\}$;
    $V_{i+1} \leftarrow V_i - M_i$;
    $S \leftarrow S \cup \{c_i\}$;
    if $V_{i+1} = \emptyset$ **then**
      for $j \leftarrow 1$ **to** i **do**
        compute $C_j = \{u \in V \ : \ \tilde{c}(u, S) = c_j\}$;
      end
      return clusters $C_1, C_2, \ldots, C_i$ centered at $c_1, c_2, \ldots, c_i$;
    end
  end
  $p_{\text{curr}} \leftarrow p_{\text{curr}}/(1 + \gamma)$;
end
Error "No clustering found";

Moreover, note that the algorithm performs at most $1+\lfloor \log_{1+\gamma} \frac{1}{p_L} \rfloor$ iterations of the while loop. Therefore, by the choice of δ' and by union bound on the number of node pairs and the number of iterations, we have that the following holds with probability $> 1 - \delta$: in each iteration every node connected to some center with probability $\geq p_{\text{curr}}$ is added to a cluster, and no cluster contains nodes whose connection probability with the center is $< (1 - \varepsilon)p_{\text{curr}}$. Note that some node with connection probability $< p_{\text{curr}}$ and $\geq (1 - \varepsilon)p_{\text{curr}}$ to a center may be added to a cluster. This implies that $p_{\text{alg}} \geq (1 - \varepsilon)p_{\text{curr}}$, with probability $\geq 1 - \delta$.

Consider now the ℓ-th iteration, with $\ell = \lfloor 2 \log_{1+\gamma}(1/p_{\text{opt}}(\tau)) \rfloor + 1$, in which we have $p_{\text{curr}} \leq p^2_{\text{opt}}(\tau)$ and the algorithm completes and returns a clustering. Note that at the beginning of this iteration we have $p_{\text{curr}} > p^2_{\text{opt}}(\tau)/(1 + \gamma)$. Therefore, by the discussion above, we have

$$p_{\text{alg}} \geq (1 - \varepsilon)p_{\text{curr}} \geq \frac{(1 - \varepsilon)}{(1 + \gamma)} p^2_{\text{opt}}(\tau)$$

with probability $\geq 1 - \delta$. By the choice of $\delta \in O\left(n^{-d}\right)$ for some $d > 1$, the statement holds with high probability. □

Consider now Algorithm 2. Algorithm 5 shows the implementation of this clustering strategy augmented with the use of sampling for estimating connection probability. Note that in this case, the threshold probability q is fixed by the user, therefore there is no need to adapt the number of samples dynamically. Therefore, we use a fixed number r of samples given by the following expression.

$$r = \frac{12}{q\varepsilon^2} \ln \frac{2n^2}{\delta} \tag{9}$$

with $\delta \in O\left(n^{-d}\right)$, for some constant $d > 1$. In the implementation, a vertex u is added to a cluster C_i of center c_i if $\tilde{q}(c_i, u) \geq q(1 - \frac{\varepsilon}{2})$, where q is the minimum connection probability threshold given in input, where estimates $\tilde{q}(c_i, u)$ are obtained through Equation (4) using the above number of samples.

Theorem 10. *With high probability, Algorithm 5 returns a partition of V into $\tau_{\text{alg}} \leq \tau(\sqrt{q})$ clusters such that each node is connected to the center of its cluster with probability at least $(1 - \varepsilon)q$.*

Algorithm 5: SEQUENTIALCLUSTER2IMPL($\mathcal{G}, q$)

$V_1 \leftarrow V$; $i \leftarrow 1$; $r \leftarrow \frac{12}{q\varepsilon^2} \ln \frac{2n^2}{\delta}$;
Sample r possible worlds $G_1, \ldots, G_r \sqsubseteq \mathcal{G}$;
while *true* **do**
  select one arbitrary cluster center $c_i \in V_i$;
  for $u \in V$ **do**
    $\tilde{q}(c_i, u) \leftarrow \frac{1}{r} \sum_{h=1}^{r} \mathbf{I}_{G_h}(u, v)$;
  end
  $C_i \leftarrow \{u \in V_i \ : \ \tilde{q}(c_i, u) \geq (1 - \frac{\varepsilon}{2})q\}$;
  $V_{i+1} \leftarrow V_i \setminus C_i$;
  if $V_{i+1} = \emptyset$ **then return** $C_1, C_2, \ldots, C_i$;
  $i \leftarrow i + 1$;
end

Proof. We first prove that each node is connected to the center of its cluster with probability at least $(1 - \varepsilon)q$. By reasoning analogously to the proof of Theorem 9, we have that with probability $1 - \delta$ no cluster contains nodes with a connection probability to the center $< (1 - \varepsilon)q$. Therefore, we have that clusters contain only nodes with connection probability to the center $\geq (1 - \varepsilon)q$, with probability $1 - \delta$.

As we did in Theorem 4, we consider an optimal $\tau(\sqrt{q})$-clustering, where each node is connected to the center of its cluster with probability at least $\sqrt{q}$. At the end of the i-th iteration of the while loop of Algorithm 5, all the nodes u that in the optimal $(\sqrt{\tau})$-clustering belong to the same cluster as one of the c_j's, with $j \leq i$, must have already been included in some cluster, since $\Pr(c_j \sim u) \geq q$ by Theorem 1. Hence, no more than $\tau_{\text{opt}}(\sqrt{q})$ clusters are returned by the algorithm. □

4.2 Parallel growing strategy

In Algorithm 3, the most critical part in a given Phase i is the computation of probability p_i, that is, the largest probability such that at least $|V_i|/2$ nodes of V_i are connected to S_i with at least this probability. Although the algorithm does not enforce a limit on the minimum probability that needs to be estimated in the phase, we can still apply a progressive sampling technique.

The implementation of our parallel cluster growing strategy with progressive sampling is provided by Algorithm 6. Let $\varepsilon \in (0, 1/2), \delta \in (0, 1)$ be suitable parameters regulating the progressive sampling, and let p_{curr} denote a probability threshold initially set to $1/(1+\varepsilon/2)$. We will use p_{curr} to bound from below the minimum probability that the algorithm will consider when adding a node to a cluster. Also, as in Algorithm 4, we need a lower bound p_L on the lowest possible probability that the algorithm may need to estimate, which can be set as explained before. For an arbitrary Phase i, let V_i represent the set of nodes that are still uncovered, S_i the set of newly selected centers, and S the set of all centers selected so far (including those in S_i). Phase i proceeds as follows. A number r of possible worlds are sampled at random, where

$$r = \frac{12}{p_{\text{curr}}\varepsilon^2} \ln \left(\frac{2n^2}{\delta} \left(2 + \left\lfloor \log \frac{1}{p_L} \right\rfloor + \lceil \log n \rceil \right) \right) \tag{10}$$

Using these r samples, the algorithm estimates the connection probabilities of all nodes of V_i with respect to all centers in S (hence, also with respect to all centers in S_i). Then, if the number of nodes $u \in V_i$ such that $\tilde{\pi}(u, S_i) \geq p_{\text{curr}} \cdot (1 + \frac{\varepsilon}{2})$ is less than $|V_i|/2$, the algorithm halves p_{curr} and repeats the estimation, increasing the number of samples accordingly. This progressive sampling is iterated until at least $|V_i|/2$ nodes satisfy the above condition. At this point, the algorithm sets p_i to the connection probability $\tilde{\pi}(u_{|V_i|/2}, S_i)$, where $u_{|V_i|/2}$ is the $|V_i|/2$-th node of V_i in a ranking based on nonincreasing values of $\tilde{\pi}(\cdot, S_i)$, and assigns each node $u \in V_i$ with $\tilde{\pi}(u, S) \geq p_i$ to the cluster centered at $\tilde{c}(u, S)$. In case the threshold probability p_{curr} drops below p_L, the algorithm aborts without producing any clustering.

Let p_{alg} be the objective function of the clustering returned by Algorithm 6 when run with arbitrary $\varepsilon \in (0, 1/2)$ and $\delta \in O\left(n^{-d}\right)$, for some constant $d > 1$. We have:

Theorem 11. *Suppose $p^2_{opt}(\tau) > 4p_L$. Algorithm 6 returns an ℓ-clustering with $\ell = O\left(\tau \log^2 n\right)$ and such that $p_{alg} \geq (1 - \varepsilon)p^2_{opt}(\tau)$, with high probability.*

Algorithm 6: ParallelClusterImpl$(\mathcal{G}, \tau)$

$V_1 \leftarrow V$; $i \leftarrow 1$; $S \leftarrow \emptyset$;
$p_{\text{curr}} \leftarrow 1/(1+\varepsilon/2)$;
while $|V_i| > \alpha\tau \log n$ **do**
  $S_i \leftarrow \{\alpha\tau \log n$ random centers from $V_i\}$;
  $S \leftarrow S \cup S_i$;
  while *true* **do**
    $r \leftarrow \frac{12}{p_{\text{curr}}\varepsilon^2} \ln\left(\frac{2n^2}{\delta}\left(2 + \left\lfloor \log \frac{1}{p_L} \right\rfloor + \lceil \log n \rceil\right)\right)$;
    Sample r possible worlds $G_1, \ldots, G_r \sqsubseteq \mathcal{G}$;
    for $c \in S$, $u \in V_i$ **do**
      $\tilde{q}(c,u) \leftarrow \frac{1}{r}\sum_{j=1}^{r} \mathbf{I}_{G_j}(c,u)$
    end
    let $u_1, u_2, \ldots, u_{V_i}$ be the nodes of V_i in decreasing order of $\tilde{\pi}(\cdot, S_i)$;
    if $\tilde{\pi}(u_{|V_i|/2}, S_i) \geq p_{curr} \cdot (1 + \frac{\varepsilon}{2})$ **then**
      $p_i \leftarrow \tilde{\pi}(u_{|V_i|/2}, S_i)$;
      foreach $u \in V_i$ **do**
        if $\tilde{\pi}(u, S) \geq p_i$ **then** assign u to cluster centered at $\tilde{c}(u, S)$;
      end
      break;
    else
      $p_{\text{curr}} \leftarrow p_{\text{curr}}/2$;
      if $p_{curr} < p_L$ **then return** "No clustering found";
    end
  end
  $V_{i+1} \leftarrow \{u \in V_i \ : \ \tilde{\pi}(u, S) < p_i\}$;
  $i \leftarrow i + 1$;
end
return the clusters centered at $S \cup V_i$;

Proof. We prove the theorem in two steps: first we will prove that if $p^2_{\text{opt}}(\tau) > 4p_L$, then the algorithm will find a clustering, with high probability; then we will show that, in each phase, there is no node v assigned to any cluster centered in c such that $\Pr(c \sim v) < (1-\varepsilon)p^2_{\text{opt}}$, with high probability. Let

$$\delta' = \frac{\delta}{n^2\left(2 + \left\lfloor \log \frac{1}{p_L} \right\rfloor + \lceil \log n \rceil\right)}$$

Assume that $p^2_{\text{opt}}(\tau) > 4p_L$, and consider an iteration of the inner while loop such that $p_{\text{curr}}(1+\varepsilon/2) \leq (1-\varepsilon/2)p^2_{\text{opt}}(\tau)$. By the choice of the number of samples and by Inequality (6), for a fixed pair of nodes u, v such that $\Pr(u \sim v) \geq p^2_{\text{opt}}(\tau)$, we have that $\tilde{q}(u,v) < (1-\varepsilon/2)p^2_{\text{opt}}(\tau)$ with probability $\leq \delta'$. By the union bound on the number of pairs, we have that $\tilde{q}(u,v) \geq (1-\varepsilon/2)p^2_{\text{opt}}(\tau) \geq (1+\varepsilon/2)p_{\text{curr}}$ for any $u, v \in V$ such that $\Pr(u \sim v) \geq p^2_{\text{opt}}(\tau)$, with probability $\geq 1-\delta$. By reasoning as in the proof of Theorem 5 we can prove that there are $|V_i|/2$ such nodes, and therefore the algorithm computes a clustering.

Consider now Phase i. We are going to prove that, for $v \in V_i$ and $c \in S$ such that $\Pr(v \sim c) < (1-\varepsilon)p^2_{\text{opt}}(\tau)$, v is added to the cluster centered in c with probability $\leq \delta'$. We condition on the event that at least one center of S_i falls into every *large cluster* as defined in Theorem 5, which happens with high probability. Let $w_1, \ldots, w_j, \ldots, w_{|V_i|/2}$ be $|V_i|/2$ nodes of V_i such that $\pi(w_j, S_i) \geq p^2_{\text{opt}}(\tau)$. For v to be added to the cluster centered in c we need that $\tilde{q}(v,c) > \left(1 + \frac{\varepsilon}{2}\right) p_{\text{curr}}$. We have two cases, which we analyze in the following.

In the first case we have $\Pr(v \sim c) < p_{\text{curr}}$. Note that this implies that $\Pr(\tilde{q}(v,c) > \left(1+\frac{\varepsilon}{2}\right) p_{\text{curr}}) < \Pr(\tilde{q}(v,c) > \left(1+\frac{\varepsilon}{2}\right) \Pr(v \sim c)) \leq \frac{\delta'}{2}$ where the last inequality follows by the choice on the number r of samples and Chernoff bound.

The second case is when $\Pr(v \sim c) \geq p_{\text{curr}}$. For $1 \leq j \leq |V_i|/2$, let c_{w_j} be the center maximizing

$\Pr(w_j \sim c_{w_j})$. Note that $\Pr(v \sim c) < \Pr(w_j \sim c_{w_j})$ by hypothesis. Moreover, observe that for v to be added to the cluster centered in c, it must be $\tilde{q}(v,c) \geq \tilde{q}(w_j, c_{w_j})$, for some $j \in [1, |V_i|/2]$. We bound the probability of this event. Note that $(1-\frac{\varepsilon}{2})\Pr(w_j \sim c_{w_j}) > (1+\frac{\varepsilon}{2})\Pr(v \sim c)$. By Chernoff bound and the choice of the number of samples

$$\begin{aligned}\Pr(\tilde{q}(w_j, c_{w_j})) &\leq \left(1-\frac{\varepsilon}{2}\right)\Pr(w_j \sim c_{w_j}) \\ &\leq e^{-\frac{r\Pr(w_j \sim c_{w_j})}{3}\frac{\varepsilon^2}{4}} \leq e^{-\frac{r p_{\text{curr}}}{3}\frac{\varepsilon^2}{4}} \leq \frac{\delta'}{2}\end{aligned}$$

and similarly

$$\Pr(\tilde{q}(v,c)) \geq \left(1+\frac{\varepsilon}{2}\right)\Pr(v \sim c) \leq \frac{\delta'}{2}$$

Therefore, $\tilde{q}(v,c) \geq \tilde{q}(w_j, c_{w_j})$ with probability $\leq \delta'$, for a fixed j.

By combining the above two cases, and by union bound on the w_j's, we have that $\tilde{q}(v,c) \geq \left(1+\frac{\varepsilon}{2}\right) p_{\text{curr}}$ with probability $\leq \delta' \cdot n$. Hence, in Phase i we have that v is added to the cluster centered in c with probability $\leq \delta' \cdot n$. By applying the union bound on all nodes v, and on the $2 + \left\lfloor \log \frac{1}{p_L} \right\rfloor + \lceil \log n \rceil$ worst case iterations of the algorithm, we have that, with probability $\geq 1-\delta$, for any cluster center c there is no node v added to the cluster centered in c such that $\Pr(v \sim c) < (1-\varepsilon)p^2_{\text{opt}}(\tau)$, and the theorem follows. □

5 Experiments

To evaluate our clustering strategies, we perform experiments along three lines. First, we compare the sequential clustering strategy of Algorithm 4 with MCL [15], which, as observed in Section 1, is a reference clustering algorithm for deterministic graphs, also used for comparison in previous works on uncertain graphs[2]. The comparison was performed with respect to various quality metrics related to the connectivity of the returned clusters, and running time. In a second set of experiments, we evaluate the performance of our sequential strategy in a predictive setting. Specifically, we consider a PPI network, using our algorithm to predict so-called *protein complexes*, and validating our results against the ground truth provided by hand-curated MIPS database of protein complexes [20]. Finally, we compare the relative performance of our sequential and parallel strategies.

graph	nodes	edges
`Collins`	1622	9074
`Gavin`	1855	7669
`Krogan`	2708	7123
`DBLP`	226413	1432920

Table 1: Graphs considered in our experiments.

We implemented our algorithms in C++[3], with the Monte Carlo sampling of possible worlds in parallel using OpenMP. Considering that the sample sizes defined in Section 4 are derived to tolerate very conservative union bounds, we verified that in practice starting the progressive sampling schedule from 50 samples always yields very accurate probability estimates. Both our code and MCL were compiled with GCC 5.4.0, and run on a Linux 4.4.0 machine equipped with an Intel I7 4-core processor and 18GB of RAM. We remark that our implementation of the parallel algorithm does not leverage on all the parallelism exposed by the concurrent growth of multiple clusters, because estimating connection probabilities in parallel is already enough to employ all the parallelism provided by our machine. Each figure we report was obtained as the average over at least 5 runs.

In our experiments we employed four different graphs, whose characteristics are summarized in Table 1. Three graphs are PPI networks, with different distributions of edge probabilities: `Collins` [21], with mostly high-probability edges; `Gavin` [22], with the majority of edges having low probabilities, and the CORE network introduced in [23] (`Krogan` in the following), which has one fourth of the edges with probability greater than 0.9, and the others almost uniformly distributed between 0.27 and 0.9. As a computationally more challenging topology, we also included in the experiments a subgraph of the DBLP collaboration network, obfuscated with the algorithm presented in [5] (reviewed in Introduction), with parameters $k = 20$ and $\varepsilon = 10^{-3}$.

[2] We were not able to compare with the algorithms proposed in [14] and [12], since their code was not made available by the authors.

[3] The source code is publicly available at `https://github.com/Cecca/ugraph`

graph	# clusters	p_{min} MCL	p_{min} SEQ	ACR MCL	ACR SEQ	$AVPR$ MCL	$AVPR$ SEQ	time (ms) MCL	time (ms) SEQ
Collins	263	0.24	0.48	0.50	0.33	0.92	0.85	377	911
	300	0.28	0.53	0.56	0.37	0.93	0.89	231	906
	317	0.44	0.56	0.60	0.39	0.94	0.90	223	629
Gavin	214	0.02	0.18	0.12	0.12	0.74	0.55	525	3400
	320	0.07	0.25	0.23	0.19	0.81	0.62	272	3320
	386	0.07	0.32	0.31	0.24	0.81	0.67	250	833
Krogan	351	0.07	0.31	0.20	0.14	0.71	0.67	1,173	3997
	580	0.16	0.43	0.36	0.23	0.72	0.76	490	4311
	705	0.16	0.57	0.43	0.30	0.71	0.83	417	1761
DBLP	17,774	≈ 0	1.00	0.89	1.0	0.99	1.0	2,002,441	2,109
	32,893	≈ 0	1.00	0.95	1.0	0.99	1.0	958,595	2,083
	42,236	≈ 0	1.00	0.96	1.0	0.99	1.0	891,443	2,102

Table 2: Comparison of MCL and SEQ. The second column reports the number of clusters returned by MCL for different inflation values (1.5, 2, and 2.5) and, consequently, the target number of nodes τ for our algorithm. Values of p_{min} reported as ≈ 0 indicate that there was a pair of nodes not connected in any sampled possible world, for 50,000 samples. Running times are in milliseconds.

5.1 Comparison with MCL

The comparison between Algorithm 4 (dubbed SEQ from now on) and MCL considers three different quality metrics. The first metric is our own, that is, the minimum connection probability p_{min} between any node and its cluster's center. (For MCL, when computing this metric we consider as cluster centers the *attractor nodes* as defined in [15].) The other two metrics are the *Average Cluster Reliability* (ACR), and the *Average Vertex Pairwise Reliability* (AVPR), both introduced in in [14], which represent, respectively, the weighted average of the probability that a cluster is connected in a random possible world, and the average connection probability of pairs of nodes belonging to the same cluster. More formally, for a given τ-clustering $\{C_1, \ldots, C_\tau\}$ of $\mathcal{G} = (V, E, p)$, we have ACR= $\sum_{i=1}^{\tau} |C_i| \Pr(\text{“}C_i \text{ is connected”})/|V|$ and AVPR = $(\sum_{i=1}^{\tau} \sum_{u,v \in C_i} \Pr(u \sim v))/\sum_{i=1}^{\tau} \binom{|C_i|}{2}$.

Recall from Section 1 that the granularity of the clustering computed by MCL cannot be controlled accurately but it is influenced by the inflation parameter. For each graph in Table 1, we run MCL with inflation set to 1.5, 2.0, and 2.5, and then we run SEQ with a value of τ which matches the granularity of the clustering returned by each MCL run. The results of the experiments are summarized in Table 2, which also reports the running times. With respect to the p_{min} metric, SEQ is always better than MCL, which is not surprising given the fact that SEQ optimizes this metric. In particular, in the case of DBLP MCL finds a clustering with an almost-zero p_{min}, meaning that there is at least one pair of nodes in the same cluster with almost zero connection probability. Conversely, SEQ finds a clustering with $p_{min} = 1$, meaning that all pairs of nodes in the same cluster are connected with probability 1, as supported by Theorem 1. As for the ACR and AVPR metrics, the scores obtained by the two algorithms are comparable, with MCL performing slightly better in most cases. We remark that p_{min} is a harder worst-case metric, whereas ACR and AVPR are averages that may hide the presence of low-probability connections in the same cluster. Also, these experiments show that for some graphs, such as DBLP, a clustering strategy that ignores the possible-world semantics, as the one at the base of MCL, may provide a very poor clustering with respect to our metric. As for the running times, MCL is almost always faster with the notable exception of the biggest DBLP graph. This is due to the fact that, thanks to progressive sampling, SEQ is able to find the clustering with $p_{min} = 1$ with a small number of samples.

5.2 Clustering as a predictive task

In PPI networks, proteins can be grouped in so-called *complexes*, that is, groups of proteins that stably interact with one another. Given a PPI network, the aim is to predict protein pairs belonging to the same complex. Our specific benchmark is the Krogan graph for which the authors published a clustering with 547 clusters obtained with MCL [23, Suppl. Table 10]. We consider a ground truth derived from the MIPS

	d	TPR	FPR	FNR
	2	0.125	0.005	0.875
	3	0.148	0.011	0.852
SEQ	4	0.168	0.035	0.832
	6	0.642	0.589	0.358
	8	0.769	0.737	0.231
MCL		0.142	0.002	0.858

Table 3: Comparison of MCL and SEQ with limited path length on `Krogan` w.r.t. the MIPS ground truth.

	# clusters		p_{min}		time (s)	
τ	PAR	SEQ	PAR	SEQ	PAR	SEQ
32	71	32	0.01	0.01	299	412
64	180	62	0.05	0.02	94	155
128	488	128	0.16	0.05	55	113
256	1031	256	0.15	0.11	74	115
512	1957	512	0.50	0.23	112	143

Table 4: Comparison of PAR and SEQ on `DBLP`. Parameter τ is the batch size for PAR and the target number of clusters for SEQ. Running times are in seconds.

database [20], which is publicly available [24]. For the purpose of the evaluation, we restrict ourselves to proteins appearing in both `Krogan` and MIPS, thus obtaining a ground truth of 13,496 protein pairs. The input to the clustering algorithms is the entire `Krogan` graph. We evaluated the returned clusterings in terms of the confusion matrix. Namely, a pair of proteins in the same cluster is considered a true positive if both proteins appear in the same MIPS complex, and a false positive otherwise; similarly, a pair of proteins in different clusters is a true negative if the two proteins do not appear in the same MIPS complex, and false negative otherwise.

We run SEQ considering only d-connection probabilities (see Section 3.3) for different values of d, and by setting $\tau = 547$ so as to match the cardinality of the reference clustering from [23]. The idea behind the use of limited path length is that we want to place in the same predicted complex proteins that are connected with a high probability, but which are also topologically close. We do not report results for $d = 1$ since there is no clustering of the `Krogan` graph with $d = 1$ and $\tau = 547$. We remark that in [12] the authors applied to the same dataset a clustering strategy based on a star-cover of the graph which is reminiscent of our algorithm with $d = 1$, but their clustering has more clusters than the reference clustering by [23]. Table 3 reports the True Positive Rate (TPR), False Positive Rate (FPR) and False Negative Rate (FNR) for SEQ with different values of d, and for MCL. We observe that, for low path lengths, our algorithm is able to find a clustering with scores similar to the reference. Note that higher values of d yield fewer false negatives at the expense of an increased number of false positives. This is expected, since with higher d the algorithm may place proteins that are topologically distant in the same predicted complex. Observe that a moderate number of false positives may be a desirable feature, since these pairs can be the target of further investigation to verify unknown protein interactions. This experiment supports our intuition that considering topologically close proteins while aiming at high connection probabilities makes our algorithm comparable with state of the art solutions in this predictive setting.

5.3 Comparison of Algorithms 4 and 6

In the last set of experiments we compare our sequential and parallel cluster growing strategies (SEQ and PAR, respectively). We use DBLP as a benchmark because its large size makes it computationally challenging for the clustering task, and compare the two algorithms in terms of the number of clusters found, the quality metric p_{min}, and the running time. We ran SEQ to find τ clusters, and PAR with batch size τ, for several values of τ. In fact, the theory developed in Section 3.2 prescribes the use of a larger batch size (i.e.. $\Theta(\tau \log n)$) to ensure that the approximation guarantees hold with high probability. However, we found that in practice a smaller batch size yields clusterings of comparable quality but smaller granularity.

We observe that PAR is faster than SEQ, even though the parallelism of the platform available for the experiment is not sufficient to unleash its full potential, due to the parallel cluster growing. Nonetheless, PAR is faster than SEQ since it does not build the clustering from scratch for each guess. PAR pays this better running time with a slight increase in the number of clusters, much lower than the theretical worst-case bound. Also, this slackness allows PAR to find a clustering with a higher p_{min}, hence making it more adaptive to the input. Observe that the values of p_{min} reported here are lower than the ones in Table 2 because different clustering granularities are considered in the two cases.

6 Conclusions

We presented a number of algorithms for clustering uncertain graphs which aim at maximizing the probability for a node to be connected to its cluster's center. Unlike all previous approaches, our algorithms feature provable guarantees on the quality of the returned clustering, while affording efficient implementations and an effective control on the granularity of the clustering. Several issues remain open for future research. From a theoretical perspective, it would be interesting to prove inapproximability results akin to those known for the standard k-center problem [18]. Also, the parallel cluster growing strategy may be improved through a more careful control on the number of clusters and/or a tighter analysis.

For what concerns the practical relevance of our clustering approach for uncertain graphs, the preliminary evidence provided by the experiments reported in the paper can be strengthened by identifying other scenarios of applicability. In this respect, further investigation will be pursued on the use of our strategies for predictivity purposes in collaboration networks. Also, more extensive experimentation on the quality-performance tradeoffs attainable by exercising different combinations of the design parameters featured by our algorithms (e.g., path length, batch size) is another promising research avenue.

References

[1] S. Asthana, O.D. King, F.D. Gibbons, and F.P. Roth. Predicting protein complex membership using probabilistic network reliability. *Genome Research*, 14(4):1170–1175, 2004.

[2] E. Adar and C. Re. Managing uncertainty in social networks. *IEEE Data Engineering Bullettin*, 30(2):15–22, 2007.

[3] A. Biswas and R. Morris. ExOR: opportunistic multi-hop routing for wireless networks. *ACM SIGCOMM Computer Communication Review*, 35(4):133–144, 2005.

[4] J. Ghosh, H.Q. Ngo, S. Yoon, and C. Qiao. On a routing problem within probabilistic graphs and its application to intermittently connected networks. In *26th IEEE INFOCOM*, pages 1721–1729, 2007.

[5] P. Boldi, F. Bonchi, A. Gionis, and T. Tassa. Injecting uncertainty in graphs for identity obfuscation. *PVLDB*, 5(11):1376–1387, 2012.

[6] Leslie G. Valiant. The complexity of enumeration and reliability problems. *SICOMP*, 8(3):410–421, 1979.

[7] M.O. Ball. Computation compexity of network reliability analysis: An overview. *IEEE Transactions on Reliability*, R-35(3):230–239, 1986.

[8] R. Jin, L. Liu, B. Ding, and H. Wang. Distance-constraint reachability computation in uncertain graphs. *PVLDB*, 4(9):551–562, 2011.

[9] M. Potamias, F. Bonchi, A. Gionis, and G. Kollios. k-nearest neighbors in uncertain graphs. *PVLDB*, 3(1):997–1008, 2010.

[10] S.E. Schaeffer. Graph clustering. *Computer Science Review*, 1(1):27–64, 2007.

[11] P. Parchas, F. Gullo, D. Papadias, and F. Bonchi. Uncertain graph processing through representative instances. *ACM TODS*, 40(3):20:1—-20:39, 2015.

[12] G. Kollios, M. Potamias, and E. Terzi. Clustering large probabilistic graphs. *IEEE TKDE*, 25(2):325–336, 2013.

[13] Yu Gu, C. Gao, G. Cong, and G. Yu. Effective and efficient clustering methods for correlated probabilistic graphs. *IEEE TKDE*, 26(5):1117–1130, 2014.

[14] L. Lin, J. Ruoming, C. Aggarwal, and S. Yelong. Reliable clustering on uncertain graphs. In *12th IEEE ICDM*, pages 459–468, Dec 2012.

[15] S. van Dongen. Graph clustering via a discrete uncoupling process. 30(1):121–141, 2008.

[16] V.V. Vazirani. *Approximation Algorithms*. Springer, 2001.

[17] M. Mitzenmacher and E. Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.

[18] T.F. Gonzalez. Clustering to minimize the maximum intercluster distance. *TCS*, 38:293–306, 1985.

[19] M. Ceccarello, A. Pietracaprina, G. Pucci, and E. Upfal. Space and time efficient parallel graph decomposition, clustering, and diameter approximation. In *27th ACM SPAA*, pages 182–191, 2015.

[20] H.W. Mewes, C. Amid, R. Arnold, D. Frishman, U. Güldener, G. Mannhaupt, M. Münsterkötter, P. Pagel, N. Strack, V. Stümpflen, et al. Mips: analysis and annotation of proteins from whole genomes. *Nucleic Acids Research*, 32(suppl 1):D41–D44, 2004.

[21] S.R. Collins, P. Kemmeren, X. Zhao, J.F. Greenblatt, F. Spencer, F. C. Holstege, J. S. Weissman, and N.J. Krogan. Toward a comprehensive atlas of the physical interactome of saccharomyces cerevisiae. *Molecular & Cellular Proteomics*, 6(3):439–450, 2007.

[22] A. C. Gavin, P. Aloy, P. Grandi, R. Krause, M. Boesche, M. Marzioch, C. Rau, L. J. Jensen, S. Bastuck, B. Dümpelfeld, et al. Proteome survey reveals modularity of the yeast cell machinery. *Nature*, 440(7084):631–636, 2006.

[23] N.J. Krogan, G. Cagney, H. Yu, G. Zhong, X. Guo, A. Ignatchenko, J. Li, S. Pu, N. Datta, A.P. Tikuisis, et al. Global landscape of protein complexes in the yeast saccharomyces cerevisiae. *Nature*, 440(7084):637–643, 2006.

[24] Institute of Bioinformatics and Systems Biology. The MIPS Comprehensive Yeast Genome Database. `ftp://ftpmips.gsf.de/fungi/Saccharomycetes/CYGD/`, May 2006.